Remote parameter estimation in a quantum spin chain enhanced by local control

Jukka Kiukas,[1] Kazuya Yuasa,[2] and Daniel Burgarth[1]

[1]*Department of Mathematics, Aberystwyth University, Aberystwyth SY23 3BZ, United Kingdom*
[2]*Department of Physics, Waseda University, Tokyo 169-8555, Japan*

We study the interplay of control and parameter estimation on a quantum spin chain. A single qubit probe is attached to one end of the chain, while we wish to estimate a parameter on the other end. We find that control on the probe qubit can substantially improve the estimation performance and discover some interesting connections to quantum state transfer.

I. INTRODUCTION

Control and estimation are two sides of the same coin. Without control, estimation cannot be performed; without estimation, the parameters required for control are unknown. On one side, the statistical inference of parameters specifying partially unknown quantum systems is currently an active research topic [1–5]. In such problems, the control resources for the estimation are usually neglected. On the other side, in the active field of quantum control [6], the knowledge of the parameters of a given system is often assumed. Given the current world-wide initiatives to build high-performance quantum devices, it is no surprise that the interplay between quantum control and estimation is regaining attention [7–10]. Only by considering both does one have the chance to find an optimal performance. We consider this problem in the context where a large many-body quantum system is observed via a single qubit probe to estimate an unknown parameter specifying a part of the Hamiltonian localized far away from the probe [11–13]. Such a setting is natural in scenarios where a large system is only partially accessible to experimentalists due to limitations forced by the implementation [14]. We show that control can substantially improve estimation and reveal some interesting connections with quantum state transfer on spin chains [15,16].

II. PARAMETER ESTIMATION ENHANCED BY CONTROL

Suppose that we wish to estimate a parameter $\lambda \in \mathbb{R}$ of a large quantum system through a single probe spin in contact with the former. We are allowed to control and measure the probe, while we have no direct access to the target system. Our Hamiltonian reads $H_{\mathbf{c},\lambda}(t) = H_\lambda + c(t)H_{\text{ctrl}}$. The target parameter λ is contained in H_λ, while H_{ctrl} represents a local control on the probe, which is tuned by the control pulse $\mathbf{c} = c(t)$. More generally, we could have more control Hamiltonians; however, in this paper we only need to consider one. We then let the system evolve from a certain initial state $|\Psi_0\rangle$, e.g., the probe prepared in $\alpha|\uparrow\rangle + \beta|\downarrow\rangle$ while the target system is initialized in its ground state. We are given a probing time T, during which we control the probe, and we measure the probe in the final state to gain some information on λ. We optimize the control field $c(t)$ to enhance the precision of the estimation of λ. See Fig. 1 for a paradigmatic setup with a spin chain.

As we can only measure the probe, the precision of the estimation is ruled by the reduced density operator $\rho_{\mathbf{c},\lambda} = \frac{1}{2}(\mathbb{1} + \mathbf{u}_{\mathbf{c},\lambda} \cdot \boldsymbol{\sigma})$ of the probe, where $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ are the Pauli operators and $\mathbf{u}_{\mathbf{c},\lambda} \in \mathbb{R}^3$ is the Bloch vector depending on the target parameter λ and the control field $\mathbf{c}$. The quantum Fisher information (QFI) [3,4] for the estimation of λ is then [17]

$$F_{\mathbf{c},\lambda} = \|\partial_\lambda \mathbf{u}_{\mathbf{c},\lambda}\|^2 + \frac{(\mathbf{u}_{\mathbf{c},\lambda} \cdot \partial_\lambda \mathbf{u}_{\mathbf{c},\lambda})^2}{1 - \|\mathbf{u}_{\mathbf{c},\lambda}\|^2}. \tag{1}$$

The optimal quantum estimator is given by $E_{\mathbf{c},\lambda} = \lambda\mathbb{1} + L_{\mathbf{c},\lambda}/F_{\mathbf{c},\lambda}$ [3], where the symmetric logarithmic derivative $L_{\mathbf{c},\lambda}$ has an explicit form given in Appendix A. It satisfies $\text{tr}[\rho_{\mathbf{c},\lambda}L_{\mathbf{c},\lambda}] = 0$ and $F_{\mathbf{c},\lambda} = \text{tr}[\rho_{\mathbf{c},\lambda}L_{\mathbf{c},\lambda}^2]$, showing that the expectation value of $E_{\mathbf{c},\lambda}$ coincides with the target parameter λ, and the variance $(\Delta E_{\mathbf{c},\lambda})^2 = 1/F_{\mathbf{c},\lambda}$ saturates the Cramér-Rao lower bound on the estimation error [3,4]. Our basic idea is to tune $\mathbf{c} = c(t)$ to maximize $F_{\mathbf{c},\lambda}$, using a tailored optimal control software.

In practice the estimator $E_{\mathbf{c},\lambda}$ must be constructed with a guessed value of λ, improved iteratively if necessary. While the general idea of enhancing the QFI by control has appeared recently in different variations (see, e.g., [8–10]), we have not found an explicit implementation of a feedback-based estimation algorithm. Hence we now proceed to describe a simple one. With the probing time T fixed, the protocol begins with an initial guess λ_0, followed by the *control step*, where a pulse $\mathbf{c} = \mathbf{c}_{\max}$ maximizing $F_{\mathbf{c},\lambda_0}$ is found. In the next step, $E_{\mathbf{c}_{\max},\lambda_0}$ is measured on $S_0 \sim 1/(\varepsilon^2 F_{\mathbf{c}_{\max},\lambda_0})$ copies of the *true* state $\rho_{\mathbf{c}_{\max},\lambda_{\text{true}}}$, and the estimate is updated to $\lambda_1 = \langle E_{\mathbf{c}_{\max},\lambda_0} \rangle$ (the average of the outcomes). The process is then iterated to get successive estimates $\lambda_1, \lambda_2, \lambda_3, \ldots$, and terminated upon reaching a predetermined accuracy ε, in the sense of $|\lambda_{n+1} - \lambda_n| < \varepsilon$. One can also use subsampling to obtain the standard error, which should tend to zero as the estimates λ_n converge, providing a better termination condition. The total number of measurements $S = \sum_n S_n$ needed for the convergence then serves as a measure of the estimation resources; by comparing S with the one obtained by running the protocol without the control step, we can judge the effect of the control on the performance of the estimation.

We use this method, together with simulated measurements (Bernoulli trials) to numerically illustrate the advantage of control in the setting introduced below, where the control turns
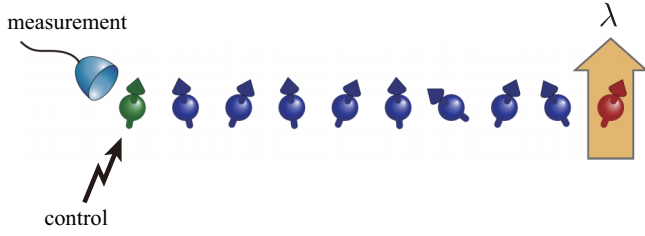
FIG. 1. A large quantum system (e.g., a chain of $N-1$ spins) accessed via a single probe spin. The probe can be measured and controlled by switching on and off local Hamiltonians. We try to enhance the precision of the estimation of an unknown parameter λ of the target system (e.g., the strength of a local magnetic field at the end spin) by the control on the probe.

out to provide a significant reduction of estimation resources for a specific nontrivial model.

III. PARAMETER ESTIMATION ACROSS A HEISENBERG CHAIN

We consider a one-dimensional chain of N spins with a Hamiltonian

$$H_{\mathbf{c},\lambda}(t) = -\frac{J}{2} \sum_{j=1}^{N-1} (\boldsymbol{\sigma}^{(j)} \cdot \boldsymbol{\sigma}^{(j+1)}) - c(t)\sigma_z^{(1)} - \lambda\sigma_z^{(N)}, \quad (2)$$

where $\boldsymbol{\sigma}^{(j)} = (\sigma_x^{(j)}, \sigma_y^{(j)}, \sigma_z^{(j)})$ are the Pauli matrices for the jth spin, $J (> 0)$ is a fixed coupling constant, $c(t)$ is a control field on the first (probe) spin (magnetic field at the first site), and λ is the unknown target parameter (magnetic field at the last site); see Fig. 1. This setting has been shown to provide full controllability over the whole chain [18], and its state transport properties have been studied in [15] without control. For quantum estimation in spin systems, see, e.g., [12,19,20].

For each spin, the eigenstates of σ_z are denoted by $|{\uparrow}\rangle$ and $|{\downarrow}\rangle$. The first spin ($j = 1$) is the probe, and we initialize it in $a|{\uparrow}\rangle + b|{\downarrow}\rangle$, with all the other spins in $|{\downarrow}\rangle$. Since $H_{\mathbf{c},\lambda}(t)$ commutes with the z component of the total spin $\sum_{j=1}^{N} \sigma_z^{(j)}$, the system remains in the subspace spanned by $|0\rangle := |{\downarrow}\rangle \otimes \cdots \otimes |{\downarrow}\rangle$, where all the spins are in $|{\downarrow}\rangle$, and $|j\rangle := |{\downarrow}\rangle \otimes \cdots \otimes |{\downarrow}\rangle \otimes |{\uparrow}\rangle \otimes |{\downarrow}\rangle \otimes \cdots \otimes |{\downarrow}\rangle$, where only the jth spin ($j = 1, \ldots, N$) is flipped to $|{\uparrow}\rangle$. The initial state of the chain is therefore $|\Psi_0\rangle = a|0\rangle + b|1\rangle$. We fix a total probing time T, and discretize it into m time slots, each of which has a different constant control. Hence the pulse $\mathbf{c} = c(t)$ is given by the vector $\mathbf{c} = (c_1, \ldots, c_m) \in \mathbb{R}^m$ determining the final state $|\Psi_{\mathbf{c},\lambda}(T)\rangle = e^{-iH_{c_m,\lambda}T/m} \cdots e^{-iH_{c_1,\lambda}T/m}|\Psi_0\rangle$. We reduce this total state to the state $\rho_{\mathbf{c},\lambda}$ of the probe spin, to which our general estimation procedure can be applied.

IV. ANALYTICAL STUDY FOR THE TWO-SPIN CASE

Before going to numerics, let us first look at the two-spin case ($N = 2$) to understand how the control helps the estimation; this is already a nontrivial indirect estimation problem. By recalling the Trotter formula, the unitary transformation induced by $H_{\mathbf{c},\lambda}(t)$ in Eq. (2) for a generic N is basically composed of the single-spin rotations on the last spin by $e^{i\lambda\sigma_z^{(N)}t}$

with the unitary transformations by the other parts of $H_{\mathbf{c},\lambda}(t)$ inserted in between. As discussed in [21], the QFI for the estimation of λ embedded by such sequential transformations is upper bounded by the optimal QFI for the estimation of λ embedded solely by $e^{i\lambda\sigma_z^{(N)}T}$ with direct initialization and measurement on the last spin permitted. That is, the QFI for our problem is upper bounded by $F_{\mathbf{c},\lambda} \leqslant 4T^2$.

If we are allowed to access all the spins, initializing and measuring them at will, this upper bound is actually achievable, by initializing and measuring the so-called NOON state $(|{\downarrow}\rangle \otimes \cdots \otimes |{\downarrow}\rangle + |{\uparrow}\rangle \otimes \cdots \otimes |{\uparrow}\rangle)/\sqrt{2}$, with no control $c(t)$ required during the probing. In our problem, however, only the probe spin is accessible, and it is not clear whether this upper bound is reachable. In particular, this puts a strong constraint on the maximal achievable QFI in the uncontrolled case. For the two-spin case, it is easy to compute the QFI in Eq. (1) for the probe spin in the absence of the control, and to perform the optimization of a and b in the initial state $|\Psi_0\rangle$. We find the asymptotic scaling of the optimized QFI for large T given by $T^2/[(1 - \lambda^2/J^2)(1 + \lambda^2/J^2)^2]$ for $\lambda^2 < 1/2$ and $4T^2(\lambda^2/J^2)/(1 + \lambda^2/J^2)^2$ otherwise. This is the best we can attain for the two-spin case without control. This QFI is upper bounded by T^2, and is smaller than the above upper bound $4T^2$ by at least a factor of 4.

The control on the probe can improve the estimation; we can even get close to the upper bound $4T^2$ if N is not large. Indeed, we provide a naive protocol consisting of three steps, and show that it is asymptotically optimal for the two-spin case. The first step is to use the control to remotely prepare a good state for sensing. The second step is to use the control to let the system evolve so as to acquire the parameter as much as possible. The final step is to remotely measure the state, by mapping the state into one that is measurable locally at the probe.

In the first step, we start with the initial state $|\Psi_0\rangle = |1\rangle = |{\uparrow}{\downarrow}\rangle$, and keep the control field at $c(t) = \lambda$ for time $\pi/(4J)$. This prepares $(|{\uparrow}{\downarrow}\rangle + i|{\downarrow}{\uparrow}\rangle)/\sqrt{2}$ up to a global phase. In the second step, we apply a very strong field $c(t) \gg J$ to suppress the exchange interaction between the two spins for time $t = T - \pi/(2J)$ [22]. Up to a global phase, the system evolves into $(|{\uparrow}{\downarrow}\rangle + ie^{2i\lambda t}|{\downarrow}{\uparrow}\rangle)/\sqrt{2}$, acquiring the relative phase depending on the target parameter λ. In the final step, we apply a strong pulse to induce an instantaneous rotation on the probe spin around the z axis to cancel the relative phase $e^{2i\lambda t}$, let the system evolve with $c(t) = \lambda$ for time $\pi/(4J)$, and measure $\sigma_z^{(1)}$ of the probe. This effectively amounts to measuring $(|{\uparrow}{\downarrow}\rangle \pm ie^{2i\lambda t}|{\downarrow}{\uparrow}\rangle)/\sqrt{2}$ in the state after the second step, which is an optimal measurement to estimate the relative phase. Note that the final state just before the measurement is $|{\downarrow}{\uparrow}\rangle$. This does not appear to depend on λ, but the value of λ taken in the first and last steps is just a guess and can be different from the true value. If the guess is not perfect, the probe fails to become the pure state $|{\downarrow}\rangle$. This failure can be detected by the measurement of $\sigma_z^{(1)}$, which helps us learn about the parameter λ. The QFI by the above procedure is calculated to be $F_{\mathbf{c},\lambda} = 4[T - (\pi/2 - 1)/J]^2$. Since the time $\pi/(2J)$ spent for the first and last steps is finite, it becomes negligible for large T, and the QFI asymptotically approaches the ultimate bound $F_{\mathbf{c},\lambda} \to 4T^2$. In this way, the control can enhance the estimation significantly.

V. NUMERICAL SIMULATION FOR A LONGER CHAIN

Let us now look at longer chains. We implement the crucial control step of the above estimation procedure using the optimal control software QTRL, which is a part of the QUTIP control package [23,24]. In addition to the existing software, we have implemented an exact gradient of the QFI in a way applicable to an arbitrary quantum system probed by a single spin. See Appendix B for details. We note that this requires the computation of the second derivatives of matrix exponentials, for which we have employed the method given in [25].

In Fig. 2, the QFI $F_{\mathbf{c},\lambda}$ normalized by T^2 with optimized control is shown as a function of the probing time T for the chain of $N = 5$ spins, and is compared with the uncontrolled case. The initial state is $|\Psi_0\rangle = |1\rangle$ for the controlled case, while it is optimized to maximize the QFI for each T for the uncontrolled case. It is clear from this result that the local control on the probe improves the precision of the estimation. In order to illustrate this explicitly, we display in the inset the results of the full estimation algorithm (described above) for the specific probing time $T = 13.5/J$: we observe that the estimation resource, i.e., the number of measurements S needed for the convergence, is significantly reduced by control, even when taking into account the considerable random variation in the results.

Longer chains naturally require longer probing times T, as ruled by the fundamental Lieb-Robinson bound for the propagation speed of the excitation, which for this specific system has been studied in [15] and found to increase linearly in N. From the estimation point of view, the relevant threshold T_0 is the time after which we start gaining the information on λ; for the $N = 5$ case, we observe from Fig. 2 that $T_0 \sim 4/J$. Another interesting time would be the one after which the Fisher information rate $F_{\mathbf{c},\lambda}/T^2$ no longer increases essentially, and the relevant question is how close to the ultimate precision bound 4 we can get. While in the two-spin case we could construct a control sequence asymptotically achieving the bound, it is now more difficult to get close to it. Loosely speaking this is because transferring the excitation to the end of the chain takes a longer time, and keeping it there to acquire the information is more difficult because the control is at the other end of the chain. Furthermore, as we see from Fig. 2, the numerics appear to become unstable for large values of T, in that different initial pulses
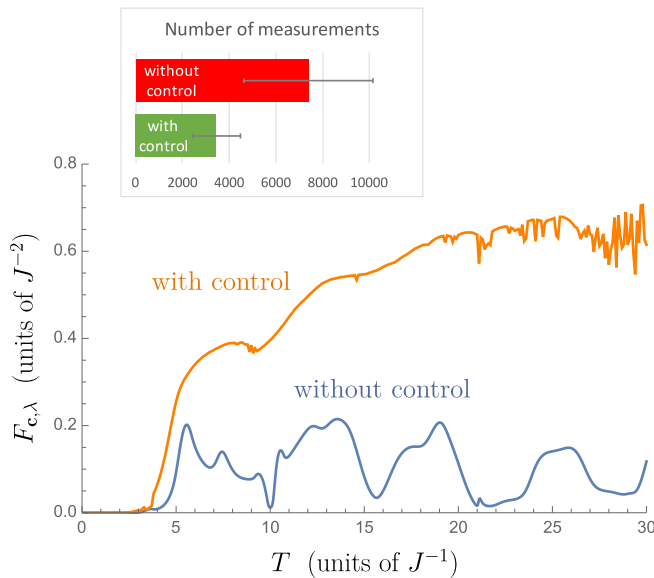

(a)

FIG. 2. Estimation of the target value $\lambda_{\text{true}} = 0$ for $N = 5$ spins. The QFI with optimized control is normalized by T^2 and shown as a function of the probing time T (orange). The initial state is chosen to be $|\Psi_0\rangle = |1\rangle$. Each probing time T is divided into 20 time slots for the control pulses $\mathbf{c}$ to be optimized. For each T, 20 numerical experiments are performed with randomly chosen initial control pulses $\mathbf{c}_0$, and the optimal QFI is taken among them. The QFI normalized by T^2 achievable without control is also shown for comparison (blue). For this uncontrolled case, the initial state $|\Psi_0\rangle = a|0\rangle + b|1\rangle$ is optimized for each T. Inset: The total number of measurements S needed to reach the target value $\lambda_{\text{true}} = 0$ with accuracy $\varepsilon = 0.01J$ starting from initial guess $\lambda_0 = 0.1J$, with and without optimal control. The probing time is $T = 13.5/J$. The value is the average over 500 numerical runs with the associated standard deviation indicated by an error bar.
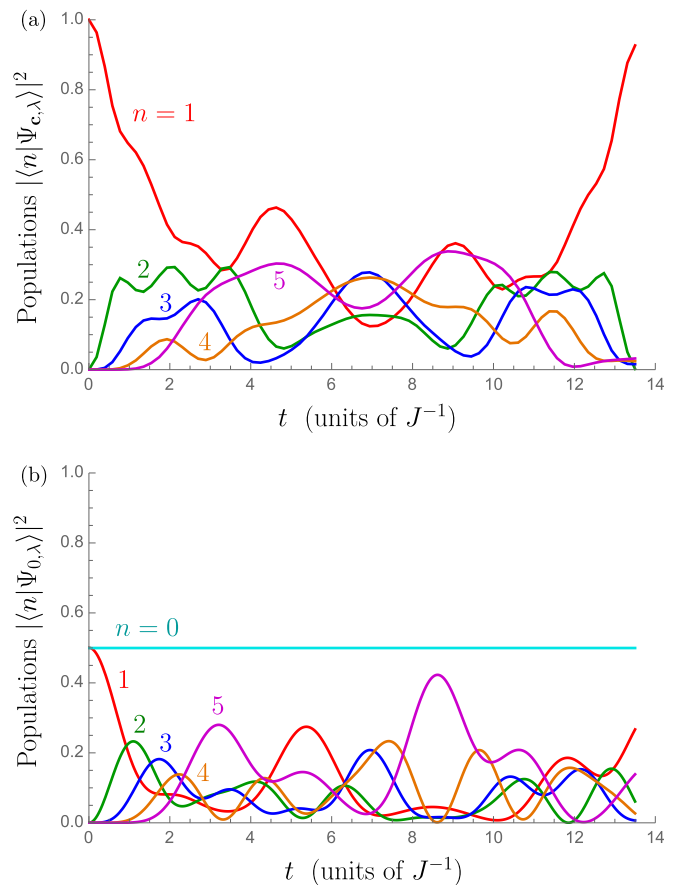
FIG. 3. Evolution of the population at each site of the chain with $N = 5$ spins for the estimation of $\lambda_{\text{true}} = 0$, (a) with an optimal control and (b) with no control, for probing time $T = 13.5/J$ divided into $m = 70$ time slots. The initial state is chosen to be $|\Psi_0\rangle = |1\rangle$ for the controlled case (a), while it is optimized to $|\Psi_0\rangle = \cos(\theta/2)|0\rangle + e^{i\varphi}\sin(\theta/2)|1\rangle$ with $(\theta,\varphi) = (1.57153, 2.41026)$ for the uncontrolled case (b). This clearly shows the structured character in the controlled case, in contrast to the uncontrolled case. In the controlled case, the excitation propagates across the chain and returns to the first site, facilitating the information transfer.

c_0 lead to different rates, making it difficult to judge if we have actually found the maximal rate. For the $N = 5$ case, we observe from Fig. 2 that after $T = 20/J$, the rate is roughly $F_{c,\lambda}/T^2 \sim 0.6$. With $N = 10, 15, 20$ we found the best rates $F_{c,\lambda}/T^2 \sim 0.1, 0.05, 0.02$, respectively ($F_{0,\lambda}/T^2 \sim 0.07, 0.03, 0.01$, respectively, in the uncontrolled case). While the QFI drops quickly with increasing the distance between the probe and the target field, the control increases it significantly.

Let us then see how the excitation is transferred across the chain. In Fig. 3, the evolution of the population probability at each site is shown for the probing time $T = 13.5/J$. In the controlled case, we start with $|\Psi_0\rangle = |1\rangle$. As in the two-spin case, we observe roughly three stages: first the excitation propagates across the chain, reaching the end at around $t \sim 2/J$. While the information on λ is accumulated, the populations remain more or less the same with some fluctuations, and eventually the excitation is brought back to the probe spin. We observed that this behavior is fairly typical also with other parameters and longer chains. In particular, the system often returns to the initial state $|1\rangle$ by the time the measurement is performed on the probe spin. As we already discussed in the two-spin case, the final state $|1\rangle$ does not explicitly depend on the target parameter λ; we learn about λ by the failure of the return of the excitation.

From the numerics it appears that the initial state $|1\rangle$ gives the best estimation results. In order to understand this, let us exclude a couple of suggestive alternatives: one might think that a superposition state of the probe such as $(|0\rangle + e^{i\phi}|1\rangle)/\sqrt{2}$ would be good for sensing the phase ϕ. Such a final pure state, as a reduced state of the probe spin, is not available for the initial state $|1\rangle$, since the probe spin is entangled with the rest of the chain when the probe spin is partially populated. It is also understandable that initial states $|\psi_0\rangle = a|0\rangle + b|1\rangle$ with small b are not useful, since the excitation transferred to the target spin is small.

VI. DISCUSSION AND OUTLOOK

We started with a natural setup for controlled parameter estimation in a quantum many-body system and found that its optimal performance is closely related to achieving perfect state transfer [15]. Although the methods developed here were used for a specific system, they are completely general. This means that we can immediately employ the program to study spin networks and experimental setups, paving the way to an optimal interplay of control and estimation in quantum technologies.

APPENDIX A: OPTIMAL ESTIMATION OF A SINGLE QUBIT

After tracing out the large system as described in the main text, we are left with the probe qubit with the density matrix

$$\rho_{c,\lambda} = \tfrac{1}{2}(\mathbb{1} + \mathbf{u}_{c,\lambda} \cdot \boldsymbol{\sigma}),$$

where $\lambda \in \mathbb{R}$ is the *unknown parameter* to be estimated, and $\mathbf{c} = c(t)$ is a control pulse to be optimized for optimal QFI. For any parameter value λ, the optimal measurement is given by

$$E_{c,\lambda} = \lambda\mathbb{1} + L_{c,\lambda}/F_{c,\lambda},$$

where the symmetric logarithmic derivative $L_{c,\lambda} = \alpha_{c,\lambda}\mathbb{1} + \mathbf{v}_{c,\lambda} \cdot \boldsymbol{\sigma}$ is determined by the coefficients

$$\alpha_{c,\lambda} = -\frac{\mathbf{u}_{c,\lambda} \cdot \partial_\lambda \mathbf{u}_{c,\lambda}}{1 - \|\mathbf{u}_{c,\lambda}\|^2},$$

$$\mathbf{v}_{c,\lambda} = \partial_\lambda \mathbf{u}_{c,\lambda} + \frac{\mathbf{u}_{c,\lambda} \cdot \partial_\lambda \mathbf{u}_{c,\lambda}}{1 - \|\mathbf{u}_{c,\lambda}\|^2}\mathbf{u}_{c,\lambda}.$$

It is easy to verify that the Hermitian operator $L_{c,\lambda}$ admits a projective measurement with two outcomes

$$\ell_{c,\lambda}^{\pm} = \alpha_{c,\lambda} \pm \|\mathbf{v}_{c,\lambda}\|$$

and projectors $(\mathbb{1} \pm \boldsymbol{\sigma} \cdot \mathbf{v}_{c,\lambda}/\|\mathbf{v}_{c,\lambda}\|)/2$.

In the estimation procedure described in the main text, the measurement of $E_{c_{\max},\lambda_n}$ with an optimal control pulse $c_{\max}$ for the nth *guessed value* λ_n is performed on the state $\rho_{c_{\max},\lambda_{\text{true}}}$, which is the *true* quantum state generated with that pulse. Hence the probabilities for the two outcomes of the measurement of $E_{c_{\max},\lambda_n}$ are given by

$$p_{\pm} = \tfrac{1}{2}(1 \pm \mathbf{u}_{c_{\max},\lambda_{\text{true}}} \cdot \mathbf{v}_{c_{\max},\lambda_n}/\|\mathbf{v}_{c_{\max},\lambda_n}\|).$$

By repeating the measurement on $S_n \sim 1/(\varepsilon^2 F_{c_{\max},\lambda_n})$ copies of the true state, we would estimate the expectation value

$$\langle E_{c_{\max},\lambda_n}\rangle = \lambda_n + (\ell_{c_{\max},\lambda_n}^+ p_+ + \ell_{c_{\max},\lambda_n}^- p_-)/F_{c_{\max},\lambda_n}. \quad (A1)$$

We simulate this experiment by generating random outcomes from a binomial distribution with sample size S_n and the true (unknown!) probabilities $p_{\pm}$ calculated from the above formula. This gives us frequencies $\tilde{p}_{\pm}$ which we then use in place of the true probabilities $p_{\pm}$ in the above formula to calculate updated estimate $\lambda_{n+1} = \langle E_{c_{\max},\lambda_n}\rangle$ described in the main text.

APPENDIX B: COMPUTATION OF THE QFI AND ITS GRADIENT

The optimal control software QRTL [24] generally optimizes control pulses by minimizing a given fidelity, which by default is chosen to be a suitable distance from a target quantum object (either a state or a unitary operator). For the purpose of the present work, we needed to implement a completely different fidelity, namely, the negative of the QFI $F_{c,\lambda}$ described in the main text. Since the optimization is based on a gradient search, the implementation requires the computation of both the function

$$c \mapsto F_{c,\lambda}$$

and its gradient. We now proceed to detail how these are obtained. We note that even though we applied the scheme to the specific Heisenberg spin chain, the composed optimization program is general and can in principle be used in any setting involving a large system estimated via an embedded single qubit probe.

First of all, we need to specify how the probe is embedded. This is done simply by defining how the qubit Pauli matrices act on the total system consisting of both the probe and the large system. For instance, if the qubit is a subsystem (i.e., a tensor factor), the action of the ith Pauli matrix is just $\tilde{\sigma}_i = \sigma_i \otimes \mathbb{1}$, where the nontrivial action is on the probe system and the identity acts everywhere else. In our case, working in a single excitation sector, it is necessary to consider direct sum instead; hence we allow arbitrary embeddings $\sigma_i \mapsto \tilde{\sigma}_i$.

As described in the main text, the final state is determined by a control pulse $\mathbf{c} = (c_1, \ldots, c_m) \in \mathbb{R}^m$, through $|\Psi_{\mathbf{c},\lambda}(T)\rangle = U_{\lambda,\mathbf{c}}|\Psi_0\rangle$, where the total transformation is a unitary $U_{\mathbf{c},\lambda} = U_{c_m,\lambda} \cdots U_{c_1,\lambda}$ given by the product of unitary propagators of the form $U_{c,\lambda} = e^{-iH_{c,\lambda}T/m}$. Here the generator $H_{c,\lambda} = H_\lambda + cH_{\text{ctrl}}$ depends on the target parameter λ and a single constant c characterizing the strength of the applied control. By tracing out everything except the probe qubit, we obtain the Bloch vector $\mathbf{u}_{\mathbf{c},\lambda}$ of the probe, which gives us the QFI $F_{\mathbf{c},\lambda}$ via Eq. (1) of the main text. The control optimization refers to maximizing $F_{\mathbf{c},\lambda}$ with respect to the pulse $\mathbf{c}$. More generally, the program can handle generators of the form $H_{c^{(1)},\ldots,c^{(M)},\lambda} = H_\lambda + \sum_{k=1}^M c^{(k)} H_{\text{ctrl}}^{(k)}$, i.e., with several control pulses, but we present the following with a single pulse so as to avoid cluttering the notation.

To compute $F_{\mathbf{c},\lambda}$, we need the components of the Bloch vector $\mathbf{u}_{\mathbf{c},\lambda}$ of the final probe state, together with their derivatives:

$$\mathbf{u}_{\mathbf{c},\lambda} = \langle U_{\mathbf{c},\lambda}\Psi_0|\tilde{\boldsymbol{\sigma}}|U_{\mathbf{c},\lambda}\Psi_0\rangle,$$

$$\partial_\lambda \mathbf{u}_{\mathbf{c},\lambda} = 2\,\text{Re}\langle \partial_\lambda U_{\mathbf{c},\lambda}\Psi_0|\tilde{\boldsymbol{\sigma}}|U_{\mathbf{c},\lambda}\Psi_0\rangle,$$

where $\tilde{\boldsymbol{\sigma}} = (\tilde{\sigma}_1, \tilde{\sigma}_2, \tilde{\sigma}_3)$ and the expectation values on the right-hand side are understood componentwise in the obvious fashion. Using the chain rule we can write the relevant derivative as

$$\partial_\lambda U_{\mathbf{c},\lambda} = \sum_{i=1}^m U_{c_m,\lambda} \cdots U_{c_{i+1},\lambda}(\partial_\lambda U_{c_i,\lambda})U_{c_{i-1},\lambda} \cdots U_{c_1,\lambda},$$

which we get once we have a method of computing $U_{c,\lambda}$ and $\partial_\lambda U_{c,\lambda}$ for arbitrary c.

We now proceed to look at the gradient. We use the shorthand notation $\partial_i = \frac{\partial}{\partial c_i}$ and suppress the parameter dependence for simplicity. By differentiating Eq. (1), we find

$$\partial_i F_{\mathbf{c},\lambda} = 2(\partial_i \partial_\lambda \mathbf{u}) \cdot \partial_\lambda \mathbf{u}$$
$$+ 2g(\partial_i \partial_\lambda \mathbf{u} \cdot \mathbf{u} + \partial_\lambda \mathbf{u} \cdot \partial_i \mathbf{u} + g\,\partial_i \mathbf{u} \cdot \mathbf{u}),$$

where $g = (\partial_\lambda \mathbf{u} \cdot \mathbf{u})/(1 - \|\mathbf{u}\|^2)$. To evaluate this function, we need two new quantities

$$\partial_i \mathbf{u} = 2\,\text{Re}\langle \partial_i U_{\mathbf{c},\lambda}\Psi_0|\tilde{\boldsymbol{\sigma}}|U_{\mathbf{c},\lambda}\Psi_0\rangle,$$

$$\partial_i \partial_\lambda \mathbf{u} = 2\,\text{Re}\langle \partial_i \partial_\lambda U_{\mathbf{c},\lambda}\Psi_0|\tilde{\boldsymbol{\sigma}}|U_{\mathbf{c},\lambda}\Psi_0\rangle$$
$$+ 2\,\text{Re}\langle \partial_\lambda U_{\mathbf{c},\lambda}\Psi_0|\tilde{\boldsymbol{\sigma}}|\partial_i U_{\mathbf{c},\lambda}\Psi_0\rangle.$$

Since c_i only appears in the ith propagator (while λ is in each one), we have

$$\partial_i U_{\mathbf{c},\lambda} = U_{c_1,\lambda} \cdots U_{c_{i-1},\lambda}(\partial_i U_{c_i,\lambda})U_{c_{i+1},\lambda} \cdots U_{c_m,\lambda},$$

$$\partial_i \partial_\lambda U_{\mathbf{c},\lambda} = (\partial_\lambda U_{c_1,\lambda}) \cdots U_{c_{i-1},\lambda}(\partial_i U_{c_i,\lambda})U_{c_{i+1},\lambda} \cdots U_{c_m,\lambda}$$
$$+ \cdots + U_{c_1,\lambda} \cdots U_{c_{i-1},\lambda}(\partial_i \partial_\lambda U_{c_i,\lambda})U_{c_{i+1},\lambda} \cdots U_{c_m,\lambda}$$
$$+ \cdots + U_{c_1,\lambda} \cdots U_{c_{i-1},\lambda}(\partial_i U_{c_i,\lambda})U_{c_{i+1},\lambda} \cdots \partial_\lambda U_{c_m,\lambda}.$$

Hence, to compute both $F_{c,\lambda}$ and its gradient, we need the four quantities $U_{c,\lambda}, \partial_i U_{c,\lambda}, \partial_\lambda U_{c,\lambda}$, and $\partial_i \partial_\lambda U_{c,\lambda}$. While the first two had already been implemented in QTRL, we needed to construct the λ derivatives specific to the estimation context.

Following [25], we observe that all these derivatives can be conveniently computed simultaneously as follows. We write $W_t = e^{-iH_{c,\lambda}t}$ for each t, and successively differentiate the evolution operator to get

$$i\partial_t W_t = H_{c,\lambda}W_t$$
$$i\partial_t\partial_\lambda W_t = (\partial_\lambda H_\lambda)W_t + H_{c,\lambda}\partial_\lambda W_t$$
$$i\partial_t\partial_c W_t = H_{\text{ctrl}}W_t + H_{c,\lambda}\partial_c W_t$$
$$i\partial_t\partial_c\partial_\lambda W_t = H_{\text{ctrl}}\partial_\lambda W_t + (\partial_\lambda H_\lambda)\partial_c W_t + H_{c,\lambda}\,\partial_c\partial_\lambda W_t.$$

This can be written as the single block matrix equation

$$i\partial_t \begin{pmatrix} W_t \\ \partial_\lambda W_t \\ \partial_c W_t \\ \partial_c\partial_\lambda W_t \end{pmatrix} = \mathbf{M}_{c,\lambda} \begin{pmatrix} W_t \\ \partial_\lambda W_t \\ \partial_c W_t \\ \partial_c\partial_\lambda W_t \end{pmatrix}$$

with

$$\begin{pmatrix} W_0 \\ \partial_\lambda W_0 \\ \partial_c W_0 \\ \partial_c\partial_\lambda W_0 \end{pmatrix} = \begin{pmatrix} \mathbb{1} \\ 0 \\ 0 \\ 0 \end{pmatrix},$$

where

$$\mathbf{M}_{c,\lambda} = \begin{pmatrix} H_{c,\lambda} & 0 & 0 & 0 \\ \partial_\lambda H_\lambda & H_{c,\lambda} & 0 & 0 \\ H_{\text{ctrl}} & 0 & H_{c,\lambda} & 0 \\ 0 & H_{\text{ctrl}} & \partial_\lambda H_\lambda & H_{c,\lambda} \end{pmatrix}.$$

Since $U_{c,\lambda} = W_{T/m}$, we immediately get

$$\begin{pmatrix} U_{c,\lambda} \\ \partial_\lambda U_{c,\lambda} \\ \partial_c U_{c,\lambda} \\ \partial_c\partial_\lambda U_{c,\lambda} \end{pmatrix} = e^{-i\mathbf{M}_{c,\lambda}T/m} \begin{pmatrix} \mathbb{1} \\ 0 \\ 0 \\ 0 \end{pmatrix},$$

that is, we can extract all the required derivatives from the first column of the blocks of the matrix $e^{-i\mathbf{M}_{c,\lambda}T/m}$. The computation of this matrix exponential therefore completes the calculation of both the QFI and its gradient.

[1] V. Giovannetti, S. Lloyd, and L. Maccone, Science **306**, 1330 (2004).

[2] M. Hayashi, *Asymptotic Theory of Quantum Statistical Inference: Selected Papers* (World Scientific, Singapore, 2005).

[3] M. G. A. Paris, Int. J. Quantum. Inform. **7**, 125 (2009).

[4] V. Giovannetti, S. Lloyd, and L. Maccone, Nat. Photon. **5**, 222 (2011).

[5] J. P. Dowling and K. P. Seshadreesan, J. Lightwave Technol. **33**, 2359 (2015).

[6] S. J. Glaser *et al.*, Eur. Phys. J. D **69**, 279 (2015).

[7] D. Burgarth and K. Yuasa, Phys. Rev. Lett. **108**, 080502 (2012).

[8] H. Yuan and C.-H. F. Fung, Phys. Rev. Lett. **115**, 110401 (2015).

[9] P. Sekatski, M. Skotiniotis, J. Kołodyński, and W. Dür, arXiv:1603.08944.

[10] J. Liu and H. Yuan, arXiv:1604.04856.

[11] C. Di Franco, M. Paternostro, and M. S. Kim, Phys. Rev. Lett. **101**, 230502 (2008).

[12] D. Burgarth, K. Maruyama, and F. Nori, Phys. Rev. A **79**, 020305 (2009).

[13] A. Sone and P. Cappellaro, Phys. Rev. A **95**, 022335 (2017).

[14] T. H. Taminiau, J. J. T. Wagenaar, T. van der Sar, F. Jelezko, V. V. Dobrovitski, and R. Hanson, Phys. Rev. Lett. **109**, 137602 (2012).

[15] S. Bose, Phys. Rev. Lett. **91**, 207901 (2003).

[16] D. Burgarth, Eur. Phys. J. ST **151**, 147 (2007).

[17] W. Zhong, Z. Sun, J. Ma, X. Wang, and F. Nori, Phys. Rev. A **87**, 022337 (2013).

[18] D. Burgarth, S. Bose, C. Bruder, and V. Giovannetti, Phys. Rev. A **79**, 060305 (2009).

[19] T. Macrì, A. Smerzi, and L. Pezzè, Phys. Rev. A **94**, 010102 (2016).

[20] F. Troiani and M. G. A. Paris, Phys. Rev. B **94**, 115422 (2016).

[21] V. Giovannetti, S. Lloyd, and L. Maccone, Phys. Rev. Lett. **96**, 010401 (2006).

[22] S. C. Benjamin and S. Bose, Phys. Rev. Lett. **90**, 247901 (2003).

[23] J. R. Johansson, P. D. Nation, and F. Nori, Comput. Phys. Commun. **184**, 1234 (2013).

[24] R. Johansson *et al.*, QUTIP/QUTIP: QUTIP-4.0.0 [Data set], Zenodo, 2016, http://doi.org/10.5281/zenodo.220867.

[25] S. Machnes, D. J. Tannor, F. K. Wilhelm, and E. Assémat, arXiv:1507.04261.